Intel Developer Update is Intel's monthly online news magazine for developers. As the official publication of developer.intel.com, it brings hardware, software, and Web developers the latest information on Intel initiatives, technologies, platforms, and products.

Cover Story

Each month, we run a cover story on the most significant industry announcement, trend, or development for the month.

Featured Articles

Delivering in-depth reports on key platforms, products and technologies, our featured articles provide a monthly source of information on issues affecting developers. Be sure to check in every month for the latest developments driving the evolution of the industry.

Contact the Editor

To make *Intel Developer Update* a better information resource, we invite you to share your thoughts on what we've published or what you'd like to see covered. Comments are always welcome.

Archives

Our archives contain two groups of previously published articles. One group contains all the articles that appeared in *Platform Solutions News*, the earlier version of *Intel Developer Update*. The articles date from September 1997 through August 1999. The other group is set up to contain *Intel Developer Update* articles dating from the inaugural September/October 1999 issue.

Bookmarking

We advise against bookmarking article pages. They're accessible online only during the month the issue is live. Thereafter, they're removed to our archives. Instead, we suggest that you bookmark the PDF (Adobe* Portable Document Format) file versions of the articles. You'll find buttons for the PDF files labeled "print article" in the right navigation section of each article. A PDF for the entire issue is labeled "print magazine" and is located near top right side of the IDU home page.

Table of Contents

(Click on page number to jump to articles)

Cover Story

Carrier-Class Voice over Packet (VoP) Architecture

Randy Smith
Senior Product Marketing Manager
Access and Switching Group
Intel Corporation

Overview

Today's networks must handle a growing variety of differentiated voice, fax, and data traffic. Voice traffic has traditionally been transferred over the dedicated circuit-switching technology of the digital Public Switched Telephone Network (PSTN). Data traffic is transferred over the lower cost, more efficient, packet-based Internet. A converged Voice over Packet (VoP) network lets carriers transport both types of traffic over a single broadband connection.

Today's implementations for high-density, carrier-class VoP processing are usually low-performance and high-cost, and they're often difficult to scale. This is because current architectures are typically based on general-purpose Digital Signal Processors (DSPs) and general-purpose packet processors that are not optimized for high-density voice processing.

Other implementations attempt to integrate voice processing and packet processing functions in a single device. However, voice processing functions are compute-intensive, while packet processing functions are control- and memory-intensive. These two distinct functions are best implemented independently in order to optimize performance in high-density VoP applications. Other issues of board real estate, power consumption, and hardware inflexibility using current solutions make it even more difficult to achieve the performance required for carrier-class VoP networks.

An optimal architecture for a VoP card should take into account the performance and scalability of digital signal processing devices, general-purpose processors, and network processors. Intel's approach is to independently optimize the architectures for the signal processor and the packet processor to achieve scalable high-channel-density solutions for carrier-class VoP applications.

Voice Processing

A voice processing card has two main types of functions:

- Control plane functions
- Data plane functions

Control plane functions use highly complex, highly differentiated software for board and device management, command interpretation, call control and signaling conversion, and messaging to call management servers.

Data plane functions are the functions most critical to the board space, power use, cost, and density of a voice processing card. Data plane functions include all (Time Division Mulitplexing) TDM-to-packet processing functions: digital signal processing, packet processing, and header processing.

Data Plane Functions

DSP voice processing functions include network echo cancellation (128-ms tail), voice compression and/or decompression, and silence suppression. DSP telephony functions include the processing of in-band signaling tones such as Dual-Tone Multi-Frequency (DTMF) and call progress tones. All these functions must be performed for numerous simultaneous active calls. They require intensive mathematical computations and lend themselves to vectorization and multiprocessing.

Packet processing functions include packetization for multiple protocols (such as RTP/UDP/IP, ATM AAL2/AAL5) and network-related Quality of Service (QoS) functions. QoS functions manage the jitter buffer, recover lost packets, generate statistics, and aggregate protocol data units from channelized DSP processing. These packet processing functions perform a range of bit manipulation operations, which are memory-intensive and control-intensive.

Design Issues

Currently, most packet processing functions are performed either by:

- An array of low-performance, high-power-consumption processors, which result in high cost
- A custom device, such as a field-programmable gate array, which results in hardware inflexibility

Neither of these solutions can deliver the performance, scalability, and flexibility required for multi-service carrier-class network equipment.

Older generation voice processing cards can use a general-purpose processor to control and manage the vectorization and multiprocessing of data plane functions. This is because older generation cards use very low-density general-purpose DSPs (1 to 12 channels). High-density next-generation cards require a different approach to controlling and managing data plane functions if they are to avoid severe system bottlenecks.

Design Approaches

There are two possible design approaches for the architecture of a high-density carrier-class voice processing card:

1. Integrate the DSP and packet processor on the same device.
2. Optimize the DSP and packet processing functions on separate devices.

The integration approach attempts to combine the math-intensive DSP functions and memory-intensive packet processor functions on the same device. However, integrating these functions requires an increase in power consumption and in the board space needed for implementation in high-channel-density systems.

Intel's approach—the optimization approach—provides a higher density, more scalable solution. This approach partitions DSP and packet processing functions onto separate devices (refer to Figure 1). In this approach, the critical data plane functions are carefully allocated to the signal processor or to the packet processor. The relationship and interactions between the processors is also carefully analyzed and defined. The result is a system-level architecture with two key elements:

- Aggregation of system memory at the most efficient level
- Distribution of control processor functionality

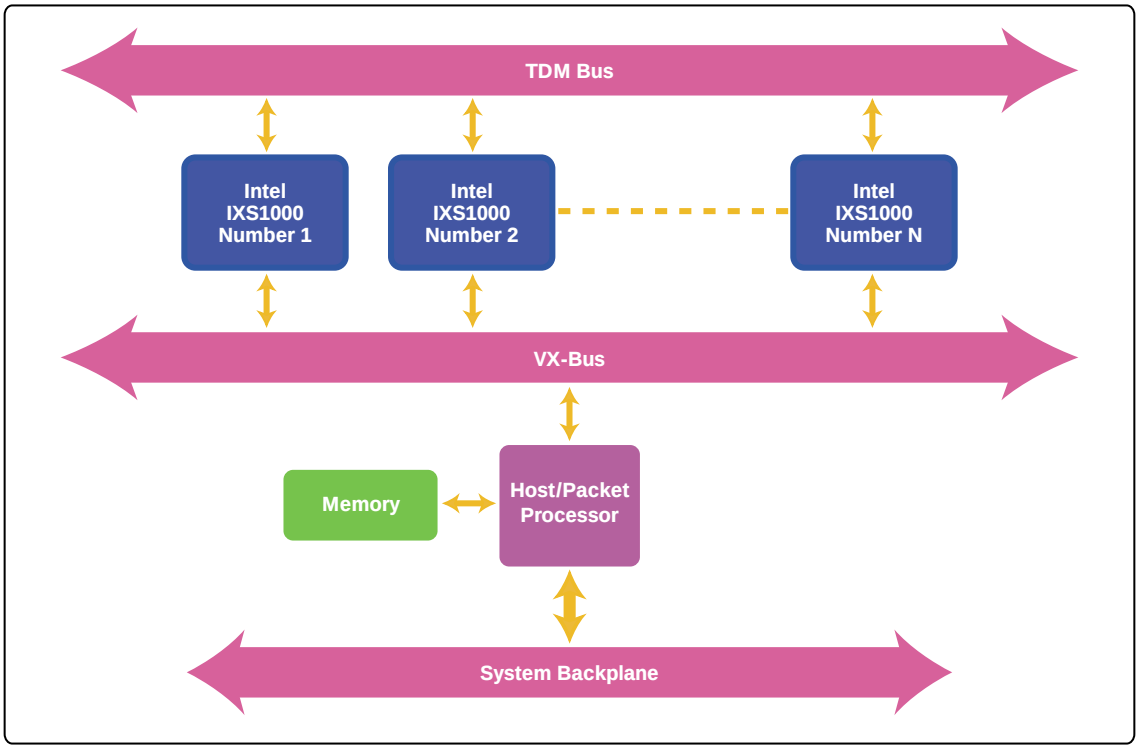

Figure 1. Voice Processing Subsystem Card based on Intel® Internet Exchange Architecture

VoP Architecture

In Intel's VoP architecture, memory is aggregated into an efficient and common external memory store, located at the packet processor. In the DSP, each integrated processing core then has dedicated local and program memory for storing and holding instructions executed within loops for signal processing.

This approach reduces by up to 50 percent the amount of memory required on the DSP. It maximizes the channels per unit area for DSP functions and eliminates the need for external memory for each DSP. Instead, the packet processor with large external memory handles the packet buffering and streaming over the packet interface (to the switch fabric or packet backplanes).

In Intel's VoP architecture, control functions are distributed within each DSP device. On a single chip, a RISC control processor is integrated with multiple DSP cores to manage and allocate DSP control tasks. The result is that the on-chip control processor performs all device control, scheduling, and flow management tasks. The DSP cores can then focus on processing data flows, enabling the device to achieve maximum performance and processing efficiency.

A third feature of Intel's VoP architecture is a hierarchical memory scheme. Hierarchical architectures reduce the memory resources needed to process each channel (lower Mbits per channel). In Intel's model, a hierarchical memory scheme makes sure that data and program segments are efficiently available for processing by the integrated DSP cores. Each DSP core has a powerful DMA engine to handle the access and transfer of programs and data in global device memory without using the valuable signal processing core resources.

Finally, Intel’s VoP architecture includes parallel processing of data segments to ensure high-performance, multi-channel signal processing. This parallelism occurs at several levels, including parallel execution on multiple cores, parallel execution of instructions on each core, single-instruction multiple-data (SIMD) for DSP tasks, and so on.

Intel® IXS1000 Media Signal Processor

In its IXS1000 Media Signal Processor, Intel offers all functional blocks necessary to interface with various network devices and buses (refer to Figure 2). This processor is a carrier-class, digital system-on-a-chip for next-generation optical networks. It integrates four high-performance signal processing cores and a control processing core with high-speed memories and intelligent I/O interfaces. It aggregates hundreds of TDM channels, performs voice/fax processing and telephony functions on each channel, then creates and transmits voice packets. The control processor manages all internal processing and I/O activities.

The IXS1000 architecture includes:

- Four digital DSP cores that execute the DSP algorithms needed to process voice and data signals
- One control processor core to schedule tasks and manage data flows for the DSPs, and to manage communication with the external host processor
- Dedicated local data memory and local program memory for each DSP core
- Four full-duplex multi-channel TDM serial interfaces
- A large global memory that stores all programs and data required for voice processing
- A high-speed internal bus and distributed DMA controllers that provide the processing cores with very fast access to data in the global memory
- A 32-bit parallel host bus interface (VX-Bus) for transferring voice packet data and for programming the device

For detailed information about the IXS1000 processor and other elements of Intel's Internet Exchange Architecture, visit the Intel Web site.

IXS1000 Block Diagram

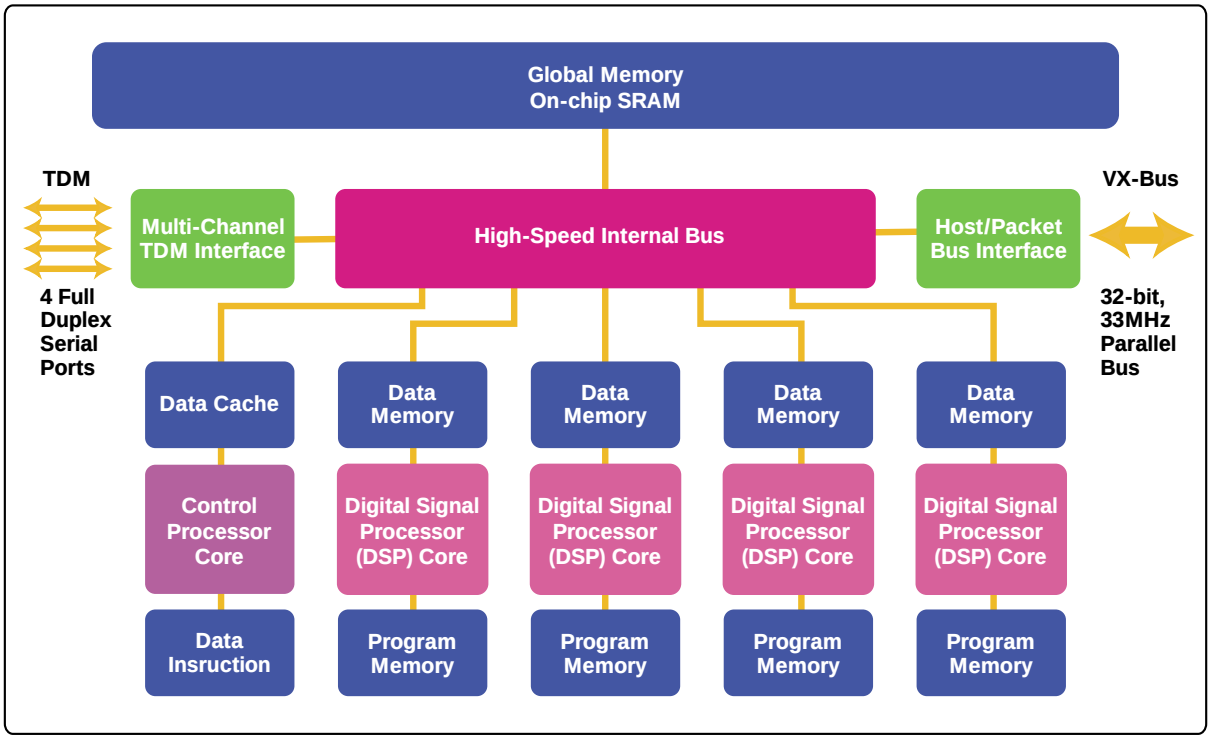

Figure 2. IXS Block Diagram

Summary

An optimal architecture for a VoP processing card optimizes the memory and compute resources of both control plane and data plane functions. An optimal, multi-service VoP solution also delivers high-density multi-channel processing in the smallest amount of space. It offers low power consumption, is programmable to add new features, and is scalable to allow for network subscriber growth.

Intel's VoP approach is specifically designed for carrier-grade voice, fax, and data processing applications. This approach segregates DSP and packet processing tasks to reduce the processing resources needed for each channel. In short, Intel's optimized VoP architecture offers higher channel density, greater flexibility, higher reliability, and lower power consumption for next-generation, carrier-class optical networks.

More Info

The IXA Voice over Packet area of the Intel Web site includes overview articles, white papers, and block diagrams about VoP technology. In particular, check out the white paper titled "Optimized High-Density Voice over Packet (VoP) Architecture."
Intel is also offering both Voice over Internet Protocol (VoIP) and Voice over ATM (VoATM) reference designs that include complete hardware, software, and application program interfaces.

You can find detailed information and specifications on the IXS1000 Signal Processor on the Intel Web site. You can also find information and specifications on the Intel® LXT3108 8-port T1/E1/J1 short- and long-haul Line Interface Unit, and the Intel® IXF3208 8-port T1/E1/J1 Framer on the site.

Author Bio

Randy Smith is a senior product marketing manager at Intel Corporation. He has over 10 years of experience in product marketing, product management, and the strategic marketing of communications and networking technologies and products with semiconductor and systems companies including SMC, AMD, and VxTel, which was recently acquired by Intel. Randy holds bachelor’s and master’s degrees in electrical engineering.

Departments

Applied Computing

The Value of Intelligence in RAID 5

Paul Luse
Senior Software Architect
I/O and Bridges Division
Intel Corporation

Overview

With the growing focus on data storage, the use of RAID technology to enhance data availability is becoming pervasive in server-attached, networked, and Internet storage applications.

Understanding how RAID works is of key importance to understanding the benefits of a particular RAID solution, including the operating characteristics and performance of particular RAID controllers and arrays.

There are many types of RAID and many various functions offered by different vendors. By understanding RAID 5, designers have a big advantage when choosing a solution. A working understanding of RAID 5 is especially important because it enables an evaluation of interrupt offload and the benefits of eXclusive OR (XOR) operations to a specific application.

If it is important to keep the server's host CPU free of the additional interrupts needed to perform a read-modify-write operation, intelligent RAID can provide a valuable solution. In addition, XOR can provide fairly significant performance boost, but it is involved only in RAID 5 writes or degraded RAID 5 reads.

Intelligent RAID

The most common RAID implementations are host-based, hardware-assisted, and intelligent RAID.

Host-based RAID (sometimes called "software RAID") runs on the host CPU, does not require special RAID hardware, and uses native drive interconnect technology. The disadvantage of host-based RAID is that the server's available application-processing bandwidth is reduced because the host CPU must devote cycles to RAID operations, including XOR calculations, data-mapping, and interrupt processing.

Hardware-assisted RAID combines a drive interconnect protocol chip with a hardware ASIC, which typically performs XOR operations. One way to think of hardware-assisted RAID is as an accelerated host-based solution. This is because the actual RAID application still executes on the host CPU, which can limit overall server performance.

With *intelligent RAID*, the host CPU is not a part of the RAID subsystem. The RAID application and XOR calculations execute on a separate I/O processor. Intelligent RAID implementations cause fewer host interrupts because they offload RAID processing from the host CPU.

RAID Levels

There are several types, or levels, of RAID, and each offers a unique set of performance and data-protection characteristics. One key concept in RAID is *abstraction*, which is the practice of "hiding" the details of an implementation to provide simplicity at a higher layer. In RAID implementations, multiple disk drives are combined into a disk array through the RAID controller and appear to the host as a single logical disk, or an "abstraction."

Here is a summary of some of the most basic and popular RAID levels:

RAID 0 is commonly referred to as "disk striping" because the RAID subsystem actually lays out the data across a number of disks in "stripes" to take advantage of parallel processing over all of the disks. A RAID 0 array is much faster for both reads and writes than a single disk because of the ability to parallel process (all disks working at the same time). RAID 0 is typically used in applications where performance requirements outweigh data-protection requirements.

RAID 1 is commonly referred to as “mirroring” because data is essentially duplicated over two or more disks. RAID 1 is typically used in applications where data protection is more important than performance. A RAID 1 array provides faster read capability than a single disk but slightly lower write performance. RAID 1 is often used to mirror the operating system boot volume of a server in a system, where protecting the operating system drive is critical.

How RAID 5 Works

RAID 5 protects the data for *n* number of disks with a single disk that is the same size as the smallest disk in the array.

For example, assume that a Web server has five disks in an application where failure of one disk must not cause server downtime. If each disk is 72 bytes, the total usable capacity for a five-disk RAID 5 array is 288 Gbytes. RAID 5 usable capacity equals *s* * [*n* - 1], where *s* is the capacity of the smallest disk in the array and *n* is the total number of disks in the array.

In this example, a single 72-Gbyte disk can guarantee that any one of the others making up the 288-Gbyte array can fail and all of the data will be safe. As another example, in a 15-disk array of 72-Gbyte disks, a single 72-Gbyte disk can protect the entire 1,008Gbyte array.

Not only does a RAID 5 array offer a very efficient way to protect data, it also has read performance similar to a RAID 0 array, while write performance suffers only slightly compared to a single disk. Because of its combination of data protection and performance, RAID 5 is very popular for general-purpose servers such as file and Web servers.

XOR Operation

How can a single disk protect the data on any number of other disks? The primary calculation is based on the very simple Boolean *XOR* operation. XOR is both an associative and commutative operation, which means that neither the order of the operands nor their grouping affects the results.

XOR is also a binary operation and only has four possible combinations of two operands. Simply put, two operands have a true XOR result when one and only one operand, exclusively, has a value of 1.

Implementing the XOR function in dedicated hardware, which can be a XOR ASIC or an I/O processor with integrated XOR functionality, greatly increases the throughput of data requiring this operation. Every byte of data stored to a RAID 5 volume requires XOR calculations. Understanding how an XOR works is critical to understanding how RAID 5 can protect so much data with such limited extra capacity.

Figure 1 represents a data map of a typical four-disk RAID 5 application.

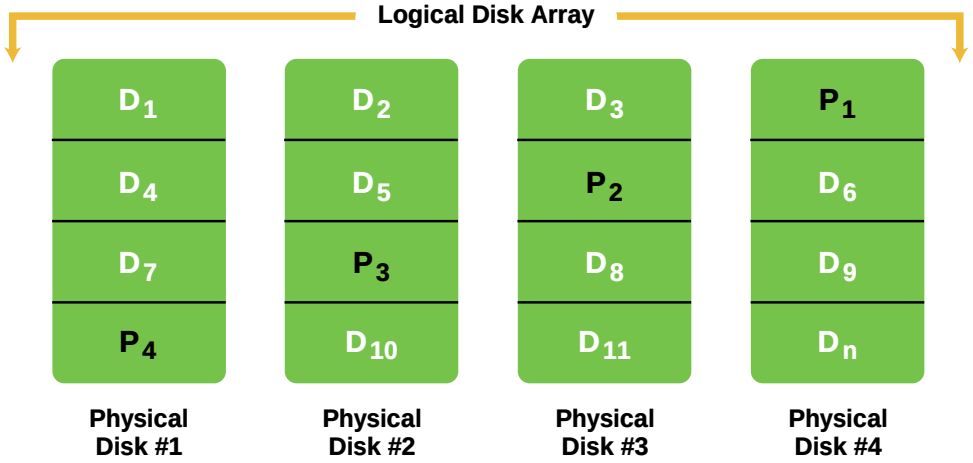

Figure 1. In RAID 5, parity data is located on a different stripe on each disk, a concept called *parity rotation.*

Note that we have introduced a new data element called P_n. P stands for *parity* data, which is simply the result of an XOR operation on all other data elements in its *stripe*. Because XOR is an associative and commutative operation, to find the XOR result of multiple operands, start by simply performing the XOR operation of any two operands. Then perform an XOR operation on the result with the next operand, and so on with all of the operands until the final result is reached.

Note that the location of the parity data is stored on a different stripe on each disk. This is called *parity rotation* and is implemented for performance reasons.

A RAID 5 volume can tolerate the loss of *any one* disk without data loss.

Figure 2 shows the array with arbitrary data values. Assume that each element represents a single bit. In real implementations, each data element would represent the total amount of data in a strip. Typical values range from 32 Kbytes to 128 Kbytes. Parity is for the first stripe: $P_1 = D_1$ XOR D_2 XOR D_3. In this example, the XOR result of D_1 and D_2 is 1, and the XOR result of 1 and D_3 is 0. Thus P_1 is 0.

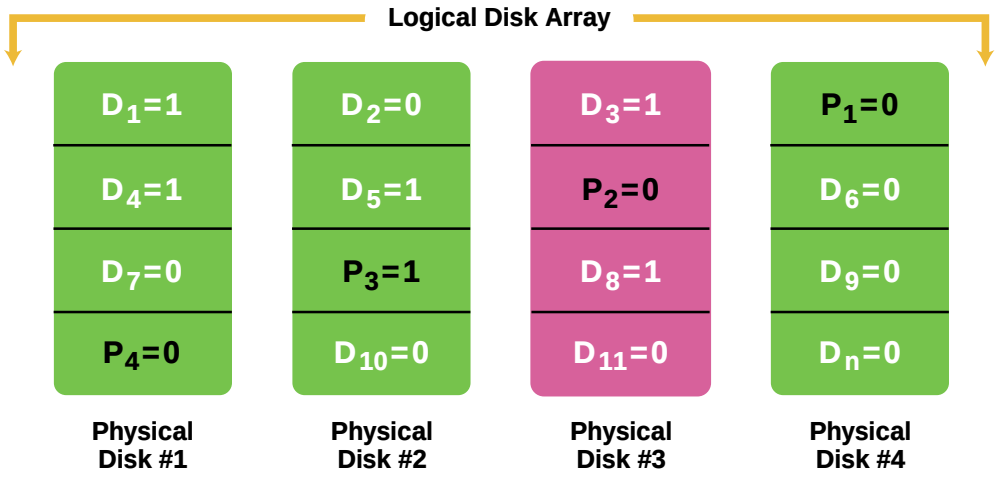

Figure 2. If a single disk fails, the missing data for any stripe is determined by performing an XOR operation on all of the remaining data elements for that stripe.

The shaded disk in Figure 2 represents a failed disk. In this situation, the disk array is typically considered *degraded.* The missing data for any stripe is easily determined by performing an XOR operation on all of the remaining data elements for that stripe.

If the host requests a RAID controller to retrieve data from a disk array that is in a degraded state, the RAID controller must first read all of the other data elements on the stripe, including the parity data element. It then performs all of the XOR calculations before it returns the data that would have resided on the failed disk. All of this happens without the host being aware of the failed disk, and array access continues. A second disk failure will result in total failure of the logical array, and the host will no longer have access.

Most RAID controllers will rebuild the array automatically if there is a spare disk available, returning the array to normal. Most RAID applications include applets or system management hooks to notify a system administrator when such a failure occurs. This notification allows the administrator to rectify the problem before another disk fails.

Read-modify-write

The RAID 5 write operation is responsible for generating parity data, which is typically referred to as a *read-modify-write* operation.

Consider a stripe composed of four strips of data and one strip of parity. Suppose the host wants to change just a small amount of data that takes up the space on only one strip within the stripe. The RAID controller cannot simply write that small portion of data and consider the request complete. It must also update the parity data. Remember that the parity data is calculated by performing XOR operations on every strip within the stripe. So when one or more strips change, parity needs to be recalculated.

Figure 3 shows a typical read-modify-write operation where the data that the host is writing to disk is contained within just one strip and identified in position D_5.

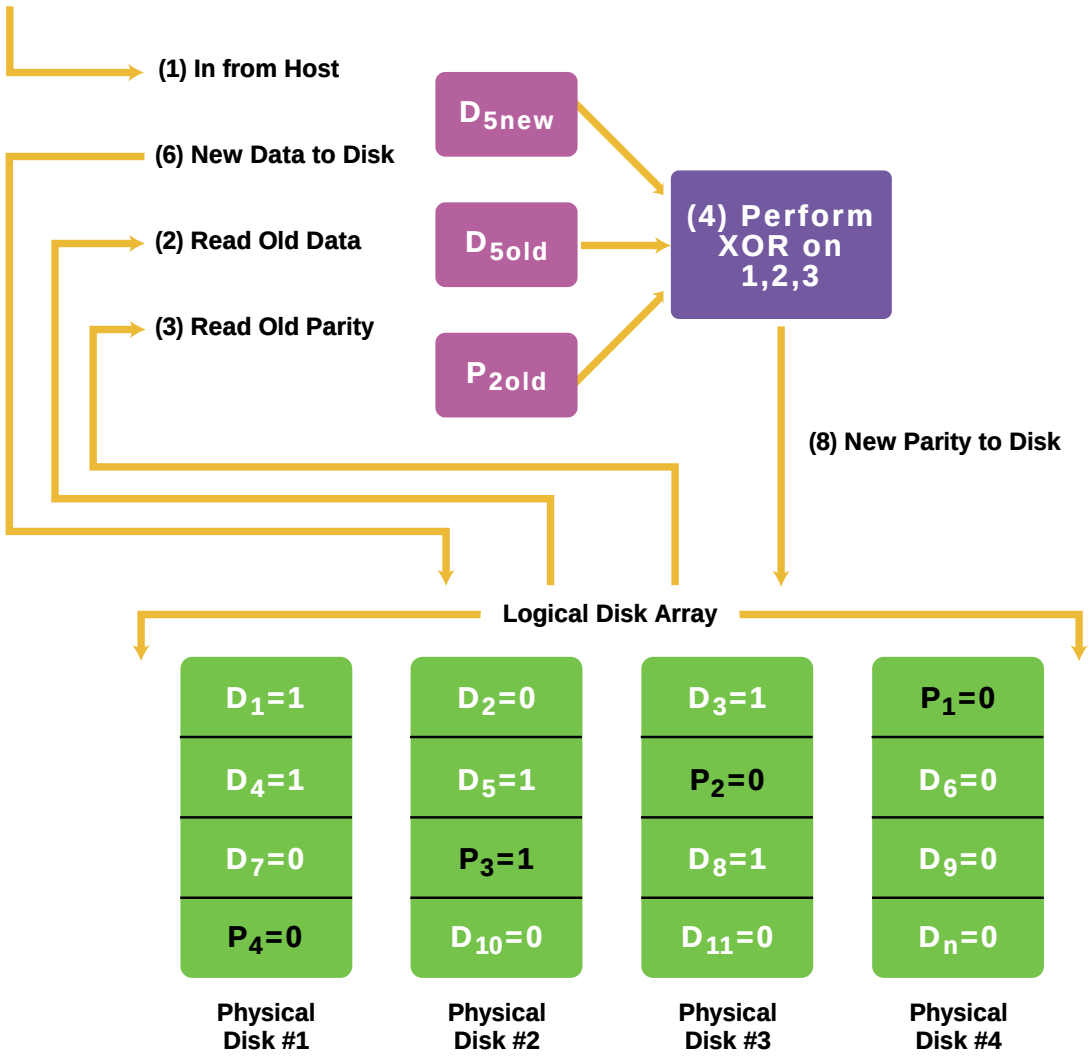

Figure 3. Steps involved in a four-disk RAID 5 read-modify-write operation.

The read-modify-write operation consists of the following steps:

1. *Read new data from host*: The host operating system requests that the RAID subsystem write a piece of data to location D_5 on disk.
2. *Read old data from target disk for new data*: Reading only the data in the location that is about to be written to eliminates the need to read all of the other disks. In addition, the number of operations involved in the read-modify-write is the same regardless of the number of disks in the array.
3. *Read old parity from target stripe for new data*: From the previous step we have the old data sitting in memory. We now pull in the old parity. This read operation is independent of the number of physical disks in the array.

4. *Calculate new parity with an XOR calculation on the data from steps 1, 2, and 3*: The XOR calculation of steps 2 and 3 provides the resultant parity of the stripe if it were totally absent of the target data element's contribution. Calculating the new parity for the stripe containing the new data simply involves performing the XOR calculation on the new data with the result of the XOR procedure performed in steps 2 and 3.
5. *Handle coherency issue*: This step is not documented in Figure 3 because its implementation varies greatly from vendor to vendor. Ensuring coherency essentially means dealing with the time period from the start of step 6 to the end of step 7. For the disk array to be considered *coherent,* or *clean,* the subsystem must ensure that the parity data block is always current for the data on the stripe. Since is not possible to guarantee that the new target data and new parity can be written to separate disks at exactly the same instant, the RAID subsystem must identify that the stripe being processed is *inconsistent*, or *dirty*, in RAID vernacular.
6. *Write new data to target location*: The data was received from the host and the RAID mappings determine which physical disk, and where on the disk, the data is to be written.
7. *Write new parity*: This also is straightforward: The new parity was calculated in step 4, now the RAID subsystem just writes it to disk.
8. *Handle coherency*: Once the RAID subsystem is assured that both steps 6 and 7 have been successfully completed and the data and parity are both on disk, the stripe is considered *coherent.*

In this example, assume that D_{5new}=0, D_{5old}=1 and P_{2old}=0. Processing step 4 on this data yields 0 $_{XOR}$ 1 $_{XOR}$ 0 = 1. This is the resultant parity element P_{2new}. The second row in Figure 5 following the read-modify-write procedure is D_4=1, D_5=0, P_2=1, and D_6=0.

This optimized method is fully scalable. The number of read, write, and XOR operations is independent of the number of disks in the array. Notice that the parity disk is involved in every write operation (steps 3 and 7). This is why parity is rotated to a different disk with each stripe. If the parity were all stored on the same disk all of the time, that disk could become a performance bottleneck.

Interrupt Offloading

An interrupt is simply a request from a system component for CPU time. I/O subsystems generate a host CPU interrupt upon completing an I/O transaction.

Consider a write operation to a four-disk RAID 5 array employing host-based RAID, hardware-assisted RAID, and intelligent RAID applications. Assume that we have the simplest of transactions: a one-bit write.

In the case of host-based RAID, the host is responsible for mapping the data to various disks. So the host must generate each read and write required to perform the read-modify-write operation. Adding them up, the host CPU should get four completion interrupts from the subsystem, consisting of two reads and two writes (steps 2, 3, 6, and 7 in the example).

A hardware-assisted RAID solution would also generate four completion interrupts because it is associated with only an XOR ASIC. The I/O processor in an intelligent RAID subsystem typically has the ability to "hide" the interim read-and-write operations from the host via various integrated peripherals. In an I/O processor-based subsystem, only a single completion interrupt is sent to the host. The I/O processor handles all of the others, freeing the host CPU to perform other non-RAID-related tasks.

Summary

Because of its combination of data protection and performance, RAID 5 is becoming very popular for general-purpose servers including file and Web servers.

With the working understanding of RAID 5 provided in this article, some of the features of RAID controllers and arrays should make more sense. The XOR operation offers a fairly significant performance boost, but is only involved in RAID 5 writes or degraded RAID 5 reads. XOR is not a factor in RAID 0 and 1. Caching is beyond the scope of this article. However, knowing that reads are actually an integral part of a RAID 5 write may help in evaluating the importance of read-cache-related features for a particular application.

A detailed understanding of a RAID 5 operation can provide a basis for evaluating the importance of interrupt offload in a specific application. If it is important to keep the host CPU free of the additional interrupts needed to perform a read-modify-write operation, intelligent RAID is a requirement.

There are many types of RAID and many various functions offered by different vendors. Understanding RAID 5 provides one basis for choosing a solution.

More Info

For additional information on Intel® Integrated RAID building blocks, including I/O processors, visit the Intel® Intelligent Internet Storage Building Blocks Web site.

Author Bio

Paul Luse is a senior software architect in Intel Communications Group's I/O and Bridges Division. He has worked in the RAID development area for more than three of his eight years with Intel. Previously, Paul has held various technical positions developing 8x930x microcontroller firmware for USB peripherals and designing internal Information Technology software solutions. He has four RAID-related U.S. patents pending.

Initiatives and Technologies

USB 2.0 Is out of the Starting Gate

Ruben Cadena
Technology Initiatives Manager
Intel Architecture Marketing Group
Intel Corporation

Overview

The USB 2.0 Specification is the next-generation peripheral connection for PCs, and it's available now. USB 2.0 offers additional I/O bandwidth over USB 1.1, and it broadens the range of next-generation peripherals that can be attached to the PC. It offers true plug-and-play functionality, including automatic device detection and installation. It's also 480 Mbits/sec—40 times faster than the original USB.

Peripheral and host-controller vendors are already offering USB 2.0-compliant products, including hard drives, CD-R/Ws, and PCI and notebook add-in cards. Other products, including scanners, printers, and cameras, will begin shipping later in 2001. Windows* operating system driver support, including a beta test program and pre-release drivers for early developers, is currently available, and Microsoft plans to ship Windows XP and Windows 2000 drivers for USB 2.0 at the same time as the Windows XP release.

Last fall, a new logo—the Certified Hi-Speed USB 2.0 logo—was introduced at Comdex*. This logo is a key part of the USB 2.0 initiative. The new logo, which indicates high-speed USB 2.0-compliant data transfer products, is available to vendors who sign a license agreement and whose products pass the required compliance testing. Testing information, pre-release USB 2.0 drivers, and information about host controllers, add-in cards, and other products can be found on the USB Implementers Forum (USB-IF) Web site. With the surge of USB 2.0-compliant peripherals hitting the market this summer, developers should insist that all storage and peripheral devices sport the new, high-speed USB 2.0 logo.

USB 2.0

The USB 2.0 Specification enables greater functionality for PCs, including higher resolution video conferencing cameras, next-generation scanners and printers, fast storage units, and faster broadband Internet connections. The higher bandwidth provided through USB 2.0 supports today's most demanding PC user applications, such as digital image creation and interactive gaming, where multiple high-speed peripherals will be running simultaneously.

Like USB 1.1, the USB 2.0 Specification satisfies the peripheral-interface needs of desktops, mobile systems, and other host platforms. For power-sensitive applications (such as notebook computers), USB 2.0 continues to offer power-management mechanisms. This allows aggressive management of I/O power consumption and ensures that USB 2.0 is beneficial even in low-power systems.

Logo Program

The USB 2.0 logo program recognizes product compliance with the USB 2.0 specification and gives consumers confidence in their USB 2.0 PCs and peripherals. There are two components to the Certified Hi-Speed USB 2.0 logo program:

- Vendors must sign a license agreement.
- Products and/or components must pass the required compliance testing.

The USB 2.0 licensing agreement, extensive test suite, and logo are available to developers, through the compliance link on the USB Web site [www.usb.org/developers]. Once vendors sign the license agreement, and their product passes the required compliance testing, that product will be listed on the Integrator's List displayed on the Web site. At that time, vendors can promote the USB 2.0 logo on their compliant products. USB 2.0-compliant products can then be listed on the USB 2.0 product search page to make it easy for retailers and buyers to locate specific USB 2.0-compliant products and components.

OS Drivers

Microsoft has stated that it plans to ship Windows XP and Windows 2000 drivers at the same time as the Windows XP release. Possible release mechanisms include OPK for OEMs and Windows Update for consumers.

Microsoft has code-complete, pre-release drivers for both Windows XP and Windows 2000. They are available for download at the USB members' Web site. To access the drivers, select the MEMBERS tab, log in, and look for USB-IF eSTORE. (Note: This area is accessible by USB-IF members only. If you are not currently a USB-IF member, then contact usmsb2@microsoft.com to request drivers directly from Microsoft.)

Microsoft has launched a USB 2.0 beta program for Windows XP and Windows 2000. To sign up for the beta program go to the USB page and follow the links and/or instructions for the beta program. The beta program runs through the end of August 2001. For non-USB-IF members, the beta-release drivers may be available through Windows Update starting in early July 2001.

Third-party vendors are also offering drivers for current and older Windows operating systems, including Windows 98, Windows 98 SE*, and Windows Me*. One such vendor is Orange Micro.

Upgrading to 2.0

The USB 2.0 Specification is fully compatible with USB 1.1 products. In fact, USB 2.0 uses the same cables and connectors as fully compliant USB 1.1 products. Developers and users can install USB 2.0 products, drivers, and/or host controllers without having to change any fully compliant USB 1.1 products or cables. This makes the transition between USB 1.1 and USB 2.0 both easy and painless.

To get the 480-Mbits/sec transfer rate, users must fully upgrade their PCs to USB 2.0 compliance. This means installing the appropriate USB 2.0 Windows driver, and installing an appropriate add-in card or discrete host controller. Without the driver and host controller (either in an add-in card or on the motherboard), the system will continue to transfer data at the original USB 1.1 rate. With the USB 2.0-compliant host controller and OS driver, the system will be enabled to transfer data at up to 40 times the original USB 1.1 rate.

Host Controllers

Currently, there are two ways to enable USB 2.0 host functions on the PC. Developers or users can install an add-in card on the PC that includes a USB 2.0-compliant host controller, or motherboard vendors can install a USB 2.0 discrete host controller on the PC motherboard.

Both add-in cards and discrete host controllers are already in production. Add-in cards are available from a variety of vendors, including Orange Micro, Adaptec, Belkin, Ratoc, and SIIG. NEC host controllers are available for PC motherboards now. Other silicon vendors will offer their host controllers later this year. Intel plans on incorporating USB 2.0 capability into its chipsets for the Intel® Pentium 4® processor beginning in the first half of 2002.

PC OEMs are also expected to introduce PCs that are USB 2.0-ready. These OEMs are working with peripheral vendors to offer bundled PCs that are fully USB 2.0-compliant later this year.

Peripherals

USB 2.0 retail peripherals are available now, and more are expected to ship this summer. Currently, storage devices from vendors such as QPS, IO Data, Melco, and others are available to the retail market. In addition, Microtek's scanner, the Plextor CD-R/W drive, and other devices—such as hard drives, CD-R/Ws, DVDs, flash card readers, scanners, printers, and digital cameras—are expected to be shipping within the next few months.

Many of these devices take advantage of USB 2.0 interoperability by connecting via USB 2.0 hubs on the PC rather than through single ports. USB 2.0 hubs are expected to start shipping in the middle of summer 2001. It's one more feature of USB 2.0 that gives users a fuller, more productive PC experience.

Event Opportunities

As USB 2.0 continues to gain momentum, there are several opportunities to get involved with USB 2.0 development activities. Near-term, the USB Implementers Forum (USB-IF) is coordinating a “USB 2.0 Exhibition area” at the Intel Developer Forum (IDF) Conference, Fall 2001 in San Jose, August 27-30. A number of silicon and peripheral vendors will be displaying their products. Additionally, there will be a USB 2.0 lab that will focus on driver performance, compliance testing, and platform design considerations. For more event information go to the USB-IF event Web site.

Summary

USB 2.0 is here now, and it’s gathering momentum. Next-generation peripheral devices for PCs are available in add-in cards, discrete host controllers, and external storage devices. The 480-Mbits/sec data-transfer speed, increased bandwidth, plug-and-play functionality, and multiple device connectivity of USB 2.0 will make USB 2.0 an exciting and ubiquitous connection for desktop and laptop PCs.

Developers can begin creating bundled PCs as well as USB 2.0-compliant products immediately. With a host of USB 2.0-compliant components and products already released, developers can participate in the evolution of USB 2.0 products. This is also an opportunity for developers to get involved with the Windows driver beta program.

More Info

The USB 2.0 specification and related information can be downloaded at no charge from the USB-IF Web site. The USB site also includes articles, vendor and OEM information, technical presentations, and information about upcoming events, including the IDF Fall ’01 Conference.

The detailed, USB technical FAQ is also available on the USB Web site. The FAQ covers specifications, compliance, signals, cables and long-haul solutions, power distribution and consumption, throughput and bandwidth, and drivers and power management.

For details on obtaining Microsoft Windows drivers, see the OS Drivers section of this article.

Author Bio

Ruben Cadena has been with Intel for nearly six years. Currently, he is a technology initiatives manager in the Intel Architecture Marketing Group at Intel. Previously, Ruben worked in electronic packaging and thermal materials.

Ruben has received several Intel Division Recognition Awards for developments in thermal technology, including his work on the "Gigamine" program. He received his B.S.E.E. from University of Texas, El Paso and his M.B.A. from Arizona State University.

PC+Consumer Electronics: the e-Home Wireless Convergence

Steve Bard
Senior Staff Architect
Intel Connected and Extended PC Lab
Intel Corporation

Overview

Glance behind any typical home entertainment center and you'll encounter a confusing "spaghetti" of interconnect cables. Replacing this cable clutter with a single high-speed, peer-to-peer digital wire improves the delivery of digital multimedia content between the devices and simplifies device configuration and control.

The consumer electronics (CE) industry has selected the High-Performance Serial Bus (IEEE 1394-1995 and its amendments, IEEE 1394a-2000 and IEEE P1394b), as the preferred interconnect between digital consumer devices. The CE industry is also developing a plan to deliver Internet content and services from an Internet broadband connection to CE devices via a 1394-connected Set-Top Box (STB). The incorporation of the Internet Protocol (IP) into 1394 makes it possible to create a seamless peer-to-peer network of CE devices and appliances known as the Consumer Entertainment Network (Ce-Net).

Within the e-Home vision, the Internet would be seamlessly available across a wide range of connected devices located throughout the home. The e-Home provides the ability to share electronic content between PCs and in-room clusters of consumer electronic devices, including digital cameras, digital camcorders, and portable music players.

Creating an e-Home environment while meeting the requirement of "no new wires" within the home presents a significant technical challenge. One solution is wireless transmission that enables the sharing of data and multimedia content between PCs and Ce-Net devices. Wireless connectivity also supports interoperability between the Ce-Net and PCs, so content traditionally reserved for delivery over a wired 1394 network can be flexibly delivered over an IEEE 802.11a/e wireless LAN (WLAN) that bridges between 1394 and 802.11.

Independent efforts have been underway within both the consumer electronics and PC industries to interconnect their respective devices with a wireless medium. To date these have been separate efforts, and they have not dealt with the issue of interoperability between CE devices and the PC.

This situation is about to change.

Wireless Convergence

We are now seeing a concerted effort to establish interoperation between the worlds of consumer electronics and personal computing. This convergence involves a wireless solution to interconnect the PC with 1394-wired clusters of CE multimedia entertainment devices, a solution that represents the missing link that can enable the PC to become a logical multimedia device with 1394 connectivity.

Wireless connectivity would make it possible for the PC to locate and identify all multimedia entertainment 1394 devices without being co-located in the same room with any specific 1394-wired cluster of devices. In addition, the PC could provide a consistent Internet interface for the connected CE devices.

The first step toward the convergence of 1394 over a wireless protocol is the creation of an industry-wide bridging standard that can support interoperability between CE devices and PCs.

1394-to-1394 Bridge over 802.11

Bridging is a term describing a mechanism that connects one method of performing functions to a different method of performing the same or similar functions. With bridging it's not possible to determine which method was used to invoke the function.

For example, the method used to transport data between PCs over Ethernet is much different from the method used to transport data between CE devices over 1394. Implementing a bridge between 1394 and Ethernet creates an environment where the PC cannot determine whether the data was originally sourced from Ethernet or from 1394.

IEEE 802.11 is a wireless standard designed to transport Ethernet datagrams used by computer networks, and it can be used as a common wireless transport to interconnect PC and CE devices. One of the principal benefits of this technology is that a 1394-to-802.11 bridge can create interoperability between CE 1394 devices and PCs. This would make it possible for the PC to be used for digital content creation or to deliver Internet content to the CE 1394 cluster.

This style of bridge does little to resolve the issue of interconnecting individual clusters of CE multimedia 1394 devices within the home. It also ignores the fact that connecting CE devices with PCs involves more than just the physical connection. The content and control methods must be interoperable.

Fortunately we can implement a wireless bridge that resolves the problem of isolated clusters of CE 1394 devices in the home and makes PCs interoperable with CE equipment. This implementation is a wired 1394-to-wired 1394 bridge over 802.11.

Bridging wired 1394-to-wired 1394 over 802.11 resolves a major issue for the CE industry, beginning with the fact that it leaves the existing infrastructure in the home untouched. There are no new wires required to connect CE multimedia entertainment 1394 clusters in various rooms of the home.

With a 1394-to-802.11 bridge software stack and an 802.11 WLAN interface card installed, the PC does not require 1394 hardware to appear as a 1394 node on the wireless infrastructure. With respect to the other 1394 devices, the PC appears to have 1394 hardware, albeit with limited functionality. In addition, the software stack enables the PC to locate and identify all of the CE multimedia 1394 devices in the 1394 clusters.

A logical PC 1394 node has limited capability when compared with that available from a physical node. Because of the real-time and bandwidth requirements, for example, digital video editing would not be possible within the constraints of a logical 1394 node. By contrast, the logical 1394 PC node can provide the ability to deliver streaming video or audio sourced from the PC, or to deliver Internet content. It can also provide command, control, and configuration applications.

When the PC and 1394 clusters are able to interact with each other, the PC will have many opportunities to provide extended services in the home.

Standard Protocol Adaptation Layer

The 1394 Trade Association Wireless Working Group (WWG) is developing a standard specification for adapting 1394 protocols and data transmission for use over a specific wireless medium: 802.11. In addition, members of the IEEE 802 organization have focused on enhancements to the Media Access Controller (MAC) layer for 802.11 (Task Group “e”) to address Quality of Service (QoS) issues.

The protocol adaptation layer (PAL) specification being developed in the 1394 Trade Association WWG is targeted specifically to work with the 802.11e MAC. The 1394 protocol and content bridges through the PAL to the 802.11e controlled wireless medium. An example of 1394 to 802.11 bridge implementation appears in Figure 1.

Bridge Implementation

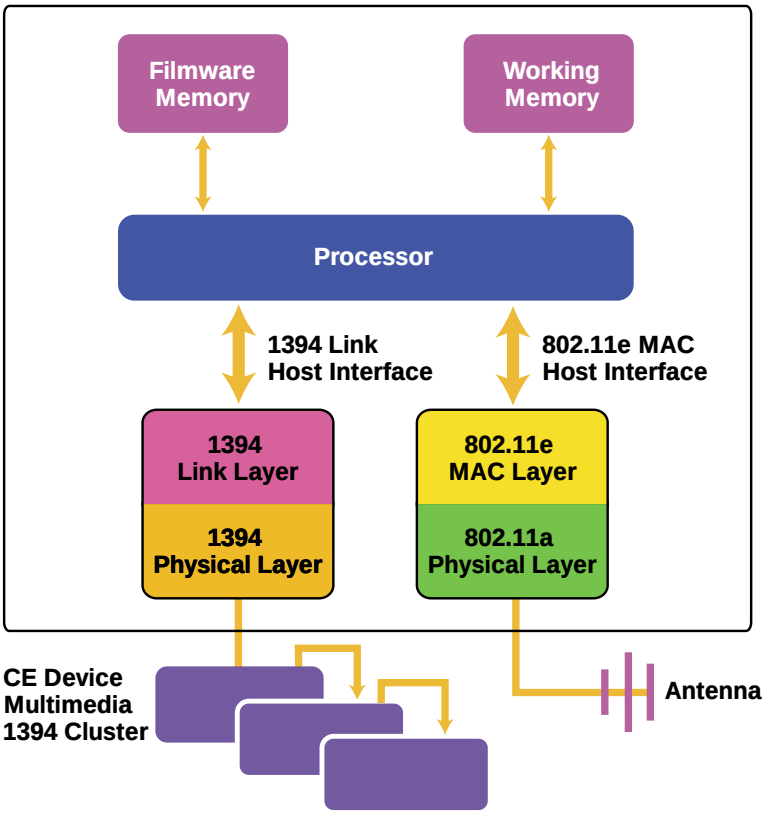

Figure 1. 1394-to-802.11e Bridge Implementation Concept

The PAL is a software/firmware transaction layer that executes on the processor within the bridge. The PAL receives 1394 content through the 1394 interface and then processes the content for delivery to the wireless medium through the 802.11e MAC.

The PAL hides the details of the 1394 layer and the 802.11e layer from each other while providing a standard set of 1394 transaction layer services to the 1394 link and a standard set of 802.11 services to the 802.11e MAC. Using these services, the PAL can emulate the high-level behavior of 1394 when presenting and accepting data to and from the 802.11e MAC while also conforming to the high-level behaviors of the wireless transport when delivering the data to its destination.

While supporting 1394 transaction layer functions (including isochronous data traffic and streaming data) the PAL must coexist with other protocols, such as IP, that use the wireless transport. To summarize: the PAL must behave like 1394, concealing the differences between the 1394 and 802.11 physical and link (MAC) layers.

User Benefits

According to leading consumer electronics manufacturers, consumers will be able to purchase 1394-equipped entertainment devices this year. Both Sony and Zenith, for example, have announced plans to market television sets equipped with 1394 interfaces during the fourth quarter of 2001. Even more consumer devices with 1394 interfaces are anticipated to become available in 2002. Within the next five to seven years, consumers will be able to purchase the majority of consumer entertainment devices with a 1394 interface. These devices will include STBs, televisions, DVD players/recorders, CD players/recorders, VHS players/recorders, AM/FM stereo receivers and amplifiers in addition to speaker systems and home theater systems.

The maze of wiring traditionally used to interconnect these devices will be replaced by a single daisy chain 1394 interconnect. Clusters of CE devices can be similarly connected in each room. The ability to deliver 1394 content over 802.11 enables the user to access content from any one room to any other room and from any device in any room to any other device, so long as the output from a device is appropriate for input to another device.

The 802.11 interface on the PC enables the PC to interchange content with any appropriate 1394 device. The PC could be used as a digital content creation tool, serving as an MP3 music repository, Internet content delivery station, or personal video recorder. The PC could also deliver a common user menu interface to every television or display device in the home. This makes it possible for the user to leverage the power of the PC to become the "conductor in the orchestra" of 1394 CE devices throughout the e-Home.

The Extended PC and R3C

Bridging the PC with CE multimedia 1394 clusters creates a robust e-Home digital environment for the consumer. It resolves the issues of how to deliver multimedia content to various clusters in the home, and how to control and configure CE devices from any in-home location.

The PC can be used to enable Remote Device Command and Configuration/Control (R3C) of every CE multimedia entertainment 1394 device in the home from any home display. In addition the PC can support an R3C application that locates and identifies all CE multimedia 1394 devices using a single universal handheld remote. This remote control device would not be too different from the television or VCR remote controls available today. This universal remote could replace all other remote control devices that exist in the home. The user could access R3C from any display device, such as a 1394-enabled TV, located anywhere in the home. Bluetooth* short-range wireless interconnect technology could be implemented in the universal handheld RF remote controller.

In this scenario, when the user presses the "select" button on the universal remote the controller activates the nearest display device. Once activated, the display device communicates via the wireless bridge to the R3C application running on the PC. The R3C application identifies the address of the target device and transmits an appropriately configured command and control GUI to that display device. The user then simply interacts with the R3C menu on the display device—selecting commands, configuring devices, and controlling the source and destination of multimedia entertainment content.

New Applications

This new level of interoperability and control opens the door to a wide range of innovative media programming options and e-Home product applications:

- The PC could provide interactive Internet Web browsing by overlaying transparent screens over the multimedia entertainment content on a TV. This type of application could be used to integrate broadcast content with interactive viewer feedback over a real-time Internet connection. This could create whole new categories of interactive programming, including home-based television game show participation.
- The PC could allow the user to select the source of multimedia content, select content from the source, and then route it to the desired display device. As a result, it would no longer be necessary to view content in the same room as the source of the content.
- To enhance e-Home security, the CE 1394 devices could be controlled by a PC software application to simulate activity in the home, thereby providing the illusion of occupancy when the home is unoccupied.
- As CE multimedia 1394 cameras become part of an interactive display device, videophones could become common, with the PC delivering streaming video to any display in the home.

Industry Support

Before the development of any hardware product capable of delivering 1394 content over a wireless medium, a standard must exist to help resolve interoperability issues.

Currently, 48 companies and organizations, including Intel, support the development of a standard for bridging 1394 over 802.11 within the 1394 Trade Association WWG. This represents a focused industrywide effort to create new product development opportunities across a very diverse global consumer electronic and PC device market.

The industry is working to resolve technical issues including:

- Available bandwidth
- Reservation of bandwidth
- Preservation of assigned bandwidth
- Quality of Service (QoS)
- Content encapsulation for transmission and reception

Next Steps

Standards development is only the beginning. Much work remains to be done, including the following steps:

- *QoS functions* must be driven into the 802.11 Media Access Controller (MAC) and must meet the requirements of 1394 isochronous traffic.
- *A working prototype*, compliant to the standards, is required to validate and improve the standards.
- *Silicon* must be developed according to the final standards.
- *New software* must be designed, developed, debugged, and tested.
- *Silicon* must be integrated into products together with the software and firmware.
- *Plug fests* must be held to facilitate true interoperability among various OEM products.

Both the consumer electronics industry and the PC industry are working together to bring to fruition an infrastructure from which a whole new class of technology will be delivered to the consumer in their e-Home.

Summary

There is an opportunity for the CE industry to take advantage of the functionality and infrastructure provided by the PC while delivering new experiences to the consumer. The PC and its connectivity infrastructure are perfectly positioned to enhance the consumer experience in the e-Home by offering CE device connectivity and configuration and control.

The PC is capable of providing proxy service that can enhance the capability of CE devices. Before this can happen, a number of significant interconnect and protocol issues must be resolved. One issue is the lack of a solution to connect one room of CE devices with CE devices in another room, without adding new wiring. To meet this challenge, standards development is now well underway to enable bridging 1394 over 802.11.

The development of such a standard represents the first step toward the creation of a PC-CE wireless bridge within the e-Home. With continued broad industry support, the PC can play a central role in the e-Home, providing opportunities for new applications and services that would not be possible otherwise.

More Info

For more information on this topic see Steve Bard’s article in Intel Technology Journal, Wireless Convergence of PC and Consumer Electronics in the e-Home.

Detailed information on the 1394 specification and industry standards activities is available on the 1394 Trade Association Web site.

Visit Intel’s Home Computing site for information on PCs, products, and ideas for the home.

Author Bio

Steve Bard has been employed by Intel for over six years and is a senior staff architect for Intel's Technology and Research Connected and Extended PC Lab (CEL). Before joining Intel, Steve worked for AST Corporation as a principal engineer in their Mobile Computing Division.

Steve has been involved in the development of 1394 technology since 1997 and has served in the IEEE 1394a-2000 Working Group. He was responsible for developing the technical details for suspend and resume (low-power capability) and worked with the group to reduce the maximum voltage on the serial bus from 40 Volts DC to 30 Volts. He also served as secretary to the IEEE P1394b Working Group and was responsible for the development of standby and restore (additional power conservation mechanisms) and the unique PIL/FOP architecture. Steve currently serves as chair for the 1394 Trade Association Wireless Working Group and is secretary to the 1394/802.11 PAL project in that group.

—End of Intel Developer Update Magazine Issue 22—